

CLAIMS

What is claimed is:

1. A method, comprising:
 - receiving a first data being written to a data block on a first storage volume;
 - indicating the data block being stored on a second storage image, the indication information being associated with a first storage image which is a copy on write snapshot; and
 - writing the first data to the data block on the second storage image.
2. The method of claim 1, wherein the indicating comprises:
 - determining whether the data block is stored on the first storage image; and
 - updating the indication information to indicate the data block is stored on the second storage image, if the data block is stored on the first storage image.
3. The method of claim 2, wherein the determining and updating comprise:
 - examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image; and
 - deleting the entry associated with the data block if the entry exists.
4. The method of claim 1, further comprising:
 - acquiring a lock from a lock mechanism before indicating the data block being stored on a second storage image; and
 - releasing the lock after writing the first data to the data block on the second storage image.

5. The method of claim 4, wherein the lock mechanism is maintained independent to the first and the second storage images.
6. The method of claim 1, further comprising:
 - receiving a second data being written to the data block on a second storage volume;
 - updating the indication information to indicate the data block is stored on the second storage image;
 - replicating an existing data stored on the data block of the second storage image to the first storage image; and
 - writing the second data to the data block on the second storage image.
7. The method of claim 6, wherein the updating comprises:
 - determining whether the data block is stored on the first storage image; and
 - updating the indication information to indicate the data block is stored on the second storage image, if the data block is stored on the first storage image.
8. The method of claim 7, wherein the determining and updating comprise:
 - examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image; and
 - creating the entry associated with the data block if the entry does not exist.
9. The method of claim 6, further comprising:
 - acquiring a lock from a lock mechanism before updating the indication information;
 - and
 - releasing the lock after writing the second data.

10. The method of claim 9, wherein the lock mechanism is maintained independent to the first and the second storage images.
11. The method of claim 1, further comprising:
 - receiving a request to read from a data block on the first storage volume;
 - determining whether the data block is stored on the first storage image or on the second storage image, based on indication information associated with the first storage image;
 - reading the data block from the first storage image if the data block is stored on the first storage image; and
 - reading the data block from the second storage image if the data block is stored on the second storage image.
12. The method of claim 11, further comprising examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image.
13. The method of claim 11, further comprising:
 - acquiring a lock from a lock mechanism before determining whether the data block is stored on the first storage image or on the second storage image; and
 - releasing the lock after reading the data block from the second storage image.
14. The method of claim 13, wherein the lock mechanism is maintained independent to the first and the second storage images.

15. A machine-readable medium having executable code to cause a machine to perform a method, the method comprising:
 - receiving a first data being written to a data block on a first storage volume;
 - indicating the data block being stored on a second storage image, the indication information being associated with a first storage image which is a copy on write snapshot; and
 - writing the first data to the data block on the second storage image.
16. The machine-readable medium of claim 15, wherein the indicating comprises:
 - determining whether the data block is stored on the first storage image; and
 - updating the indication information to indicate the data block is stored on the second storage image, if the data block is stored on the first storage image.
17. The machine-readable medium of claim 16, wherein the determining and updating comprise:
 - examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image; and
 - deleting the entry associated with the data block if the entry exists.
18. The machine-readable medium of claim 15, wherein the method further comprises:
 - acquiring a lock from a lock mechanism before indicating the data block being stored on a second storage image; and
 - releasing the lock after writing the first data to the data block on the second storage image.

19. The machine-readable medium of claim 18, wherein the lock mechanism is maintained independent to the first and the second storage images.
20. The machine-readable medium of claim 15, wherein the method further comprises:
 - receiving a second data being written to the data block on a second storage volume;
 - updating the indication information to indicate the data block is stored on the second storage image;
 - replicating an existing data stored on the data block of the second storage image to the first storage image; and
 - writing the second data to the data block on the second storage image.
21. The machine-readable medium of claim 20, wherein the updating comprises:
 - determining whether the data block is stored on the first storage image; and
 - updating the indication information to indicate the data block is stored on the second storage image, if the data block is stored on the first storage image.
22. The machine-readable medium of claim 21, wherein the determining and updating comprise:
 - examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image; and
 - creating the entry associated with the data block if the entry does not exist.
23. The machine-readable medium of claim 20, wherein the method further comprises:
 - acquiring a lock from a lock mechanism before updating the indication information;
 - and

releasing the lock after writing the second data to the data block on the second storage image.

24. The machine-readable medium of claim 23, wherein the lock mechanism is maintained independent to the first and the second storage images.
25. The machine-readable medium of claim 15, wherein the method further comprises:
 - receiving a request to read from a data block on the first storage volume;
 - determining whether the data block is stored on the first storage image or on the second storage image, based on indication information associated with the first storage image;
 - reading the data block from the first storage image if the data block is stored on the first storage image; and
 - reading the data block from the second storage image if the data block is stored on the second storage image.
26. The machine-readable medium of claim 25, wherein the method further comprises examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image.
27. The machine-readable medium of claim 25, wherein the method further comprises:
 - acquiring a lock from a lock mechanism before determining whether the data block is stored on the first storage image or on the second storage image; and
 - releasing the lock after reading the data block from the second storage image.

28. The machine-readable medium of claim 27, wherein the lock mechanism is maintained independent to the first and the second storage images.
29. An apparatus, comprising:
means for receiving a first data being written to a data block on a first storage volume;
means for indicating the data block being stored on a second storage image, the indication information being associated with a first storage image which is a copy on write snapshot; and
means for writing the first data to the data block on the second storage image.
30. The apparatus of claim 29, wherein the means for indicating comprises:
means for determining whether the data block is stored on the first storage image; and
means for updating the indication information to indicate the data block is stored on the second storage image, if the data block is stored on the first storage image.
31. The apparatus of claim 30, wherein the means for determining and means for updating comprise:
means for examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image; and
means for deleting the entry associated with the data block if the entry exists.
32. The apparatus of claim 29, further comprising:
means for acquiring a lock from a lock mechanism before indicating indicating the data block being stored on a second storage image; and
means for releasing the lock after writing the first data to the data block on the second storage image.

33. The apparatus of claim 32, wherein the lock mechanism is maintained independent to the first and the second storage images.
34. The apparatus of claim 29, further comprising:
- means for receiving a second data being written to the data block on a second storage volume;
 - means for updating the indication information to indicate the data block is stored on the second storage image;
 - means for replicating an existing data stored on the data block of the second storage image to the first storage image; and
 - means for writing the second data to the data block on the second storage image.
35. The apparatus of claim 34, wherein the means for updating comprises:
- means for determining whether the data block is stored on the first storage image; and
 - means for updating the indication information to indicate the data block is stored on the second storage image, if the data block is stored on the first storage image.
36. The apparatus of claim 35, wherein the means for determining and means for updating comprise:
- means for examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image; and
 - means for creating the entry associated with the data block if the entry does not exist.
37. The apparatus of claim 34, further comprising:

means for acquiring a lock from a lock mechanism before updating the indication information; and
means for releasing the lock after writing the second data.

38. The apparatus of claim 37, wherein the lock mechanism is maintained independent to the first and the second storage images.

39. The apparatus of claim 29, further comprising:

means for receiving a request to read from a data block on the first storage volume;
means for determining whether the data block is stored on the first storage image or on a second storage image, based on indication information associated with the first storage image;
means for reading the data block from the first storage image if the data block is stored on the first storage image; and
means for reading the data block from the second storage image if the data block is stored on the second storage image.

40. The apparatus of claim 39, further comprising means for examining a lookup table whether there is an entry associated with the data block, the lookup table being associated with the first storage image.

41. The apparatus of claim 39, further comprising:

means for acquiring a lock from a lock mechanism before determining whether the data block is stored on the first storage image or on a second storage image;
and

means for releasing the lock after reading the data block from the second storage image.

42. The apparatus of claim 41, wherein the lock mechanism is maintained independent to the first and the second storage images.
43. A data storage system, comprising:
a processing system; and
a memory coupled to the processing system, the memory storing instructions, which when executed by the processing system, cause the processing system to perform the operations of:
receiving a first data being written to a data block on a first storage volume;
indicating the data block being stored on a second storage image, the indication information being associated with a first storage image which is a copy on write snapshot; and
writing the first data to the data block on the second storage image.
44. A method for preserving data in a data processing system, the method comprising:
receiving a command to preserve data in a data storage system;
executing, for a first data, a first I/O (input/output) process directed to a first storage volume, wherein the first I/O process begins at a first time which is prior to receiving the command;
creating a data structure for a copy on write snapshot in a second image which corresponds to a second storage volume; and
writing a second data directed to the second storage volume as part of a second I/O process which begins after the first time, wherein the writing of the second I/O

process writes the second data to a first image corresponding to the first storage volume and wherein data in the data structure is modified if data in the second image corresponds to the second data.

45. The method of claim 44, wherein the first I/O process is a write operation to storage.
46. The method of claim 44, wherein the first image is stored on the first storage volume and wherein the data structure is created as a result of receiving the command.
47. The method of claim 44, wherein the second I/O process is capable of executing while the first I/O process is executing.
48. The method of claim 44, wherein the second I/O process writes data to the first image without writing to the second image.
49. The method of claim 44, wherein the second I/O process copies data in the first image to the second image and creates an entry in the data structure for the data copied to the second image.
50. The method of claim 44, wherein the first storage volume is a first virtual logical unit (VLUN) and the second storage volume is a second (VLUN).
51. A machine-readable medium having executable code to cause a machine to perform a method, the method comprising:
receiving a command to preserve data in a data storage system;

executing, for a first data, a first I/O (input/output) process directed to a first storage volume, wherein the first I/O process begins at a first time which is prior to receiving the command;

creating a data structure for a copy on write snapshot in a second image which corresponds to a second storage volume; and

writing a second data directed to the second storage volume as part of a second I/O process which begins after the first time, wherein the writing of the second I/O process writes the second data to a first image corresponding to the first storage volume and wherein data in the data structure is modified if data in the second image corresponds to the second data.

52. The machine-readable medium of claim 51, wherein the first I/O process is a write operation to storage.
53. The machine-readable medium of claim 51, wherein the first image is stored on the first storage volume and wherein the data structure is created as a result of receiving the command.
54. The machine-readable medium of claim 51, wherein the second I/O process is capable of executing while the first I/O process is executing.
55. The machine-readable medium of claim 51, wherein the second I/O process writes data to the first image without writing to the second image.

56. The machine-readable medium of claim 51, wherein the second I/O process copies data in the first image to the second image and creates an entry in the data structure for the data copied to the second image.
57. The machine-readable medium of claim 51, wherein the first storage volume is a first virtual logical unit (VLUN) and the second storage volume is a second (VLUN).
58. An apparatus, comprising:
- means for receiving a command to preserve data in a data storage system;
 - means for executing, for a first data, a first I/O (input/output) process directed to a first storage volume, wherein the first I/O process begins at a first time which is prior to receiving the command;
 - means for creating a data structure for a copy on write snapshot in a second image which corresponds to a second storage volume; and
 - means for writing a second data directed to the second storage volume as part of a second I/O process which begins after the first time, wherein the writing of the second I/O process writes the second data to a first image corresponding to the first storage volume and wherein data in the data structure is modified if data in the second image corresponds to the second data.
59. The apparatus of claim 58, wherein the first I/O process is a write operation to storage.
60. The apparatus of claim 58, wherein the first image is stored on the first storage volume and wherein the data structure is created as a result of receiving the command.

61. The apparatus of claim 58, wherein the second I/O process is capable of executing while the first I/O process is executing.
62. The apparatus of claim 58, wherein the second I/O process writes data to the first image without writing to the second image.
63. The apparatus of claim 58, wherein the second I/O process copies data in the first image to the second image and creates an entry in the data structure for the data copied to the second image.
64. The apparatus of claim 58, wherein the first storage volume is a first virtual logical unit (VLUN) and the second storage volume is a second (VLUN).
65. A data storage system, comprising:
 - a processing system; and
 - a memory coupled to the processing system, the memory storing instructions, which when executed by the processing system, cause the processing system to perform the operations of:
 - receiving a command to preserve data in the data storage system;
 - executing, for a first data, a first I/O (input/output) process directed to a first storage volume, wherein the first I/O process begins at a first time which is prior to receiving the command;
 - creating a data structure for a copy on write snapshot in a second image which corresponds to a second storage volume; and
 - writing a second data directed to the second storage volume as part of a second I/O process which begins after the first time, wherein the writing of the

second I/O process writes the second data to a first image corresponding to the first storage volume and wherein data in the data structure is modified if data in the second image corresponds to the second data.

- 66. The method of claim 44, wherein the second I/O process is capable of accessing the same data, via the second storage volume, as the first I/O process.
- 67. The machine-readable medium of claim 51, wherein the second I/O process is capable of accessing the same data, via the second storage volume, as the first I/O process.
- 68. The apparatus of claim 58, wherein the second I/O process is capable of accessing the same data, via the second storage volume, as the first I/O process.